


SGI Origin 3400 Server



The SGI Origin 3400 server comes in one of four easy-to-order preconfigured bundles. Based on the tall rack, SGI Origin 3400 integrates C-bricks, R-bricks, and any of the three I/O bricks into a single rack. For maximum I/O and processor configurations, an optional second I/O-only rack can be added. Two R-bricks with 6-port routers come standard in all configurations, which allow simple upgrades to a maximum of eight C-bricks. Four C-bricks are connected to each R-brick, and the two R-bricks are tightly linked together with two NUMalink cables for maximum bandwidth.

SGI Origin 3400 Technical Specifications


A single 39U rack configuration with additional racks for I/O and disks as required:

Description	Minimum	Maximum
Memory size	512MB (one C-brick with one 512MB bank installed)	64GB (eight C-bricks, each with four 2GB banks installed)
Processors	4 processors (one C-brick)	32 processors (eight C-bricks)
Input/output	One I-brick	One I-brick, seven optional I/O bricks
R-bricks	Two 6-port R-bricks	Two 6-port R-bricks
Power bays	One power bay	Two power bays

C-brick processor configuration options:

Standard Bundles	Processor Base Configurations	Processor Upgrade Possibilities
4 processors	One 4P C-brick and 2 R-bricks	Add 1-7 4P C-bricks
8 processors	Two 4P C-bricks and 2 R-bricks	Add 1-6 4P C-bricks
16 processors	Four 4P C-bricks and 2 R-bricks	Add 1-4 4P C-bricks
32 processors	Eight 4P C-bricks and 2 R-bricks	N/A

SGI Onyx 3400 System



The SGI Onyx 3400 system comes in one of four easy-to-order preconfigured bundles. Based on the tall rack, SGI Onyx 3400 integrates C-bricks, R-bricks, and any of the three I/O bricks into the first rack, and a minimum of one G-brick in a second graphics rack. For maximum I/O and processor configurations, an optional third I/O-only rack can be added. Two R-bricks with 6-port routers come standard in all configuration, which allow simple upgrades to a maximum of eight C-bricks, all in the same first (CPU) rack. SGI Onyx 3400 supports up to 4 G-bricks and 8 graphics pipes, based on InfiniteReality2 or InfiniteReality3 graphics.

SGI Onyx 3400 Technical Specifications


A two 39-U rack configuration with additional racks for I/O and disks as required:

Description	Minimum	Maximum
Memory size	512MB (one C-brick with one 512MB bank installed)	64GB (eight C-bricks, each with four 2GB banks installed)
Processors	4 processors (one C-brick)	32 processors (eight C-bricks)
Graphics Module	One G-brick (pipe not included)	4 G-bricks
Graphics pipe	One InfiniteReality2 or IR3 pipe, one RM, one DG, one monitor	8 graphics pipes
Graphics options	Keyboard, mouse, audio card, USB extender	8 sets of options, one for each graphics pipe
Input/output	One I-brick	One I-brick, seven optional I/O bricks
R-bricks	Two 6-port R-bricks	Two 6-port R-bricks
Power bays	One power bay	Two power bays

C-brick processor configuration options:

Standard Bundles	Processor Base Configurations	Processor Upgrade Possibilities
4 processor	One 4P C-brick and 2 R-bricks	Add 1-7 4P C-bricks
8 processor	Two 4P C-bricks and 2 R-bricks	Add 1-6 4P C-bricks
16 processor	Four 4P C-bricks and 2 R-bricks	Add 1-4 4P C-bricks
32 processor	Eight 4P C-bricks and 2 R-bricks	N/A

SGI Origin 3800 Server



The SGI Origin 3800 server entry configuration includes two R-bricks and four C-bricks in the compute rack and one I-brick and one P-brick in the second (I/O) rack. The two R-bricks have 8-port routers, which are standard in all SGI Origin 3800 configurations. Upgrades can be added in single C-brick increments or by easy-to-order 32-processor full racks. Configurations include more than 128 processors require an additional R-brick called the Metarouter, along with extra NUMALink cables, to scale the system up to a maximum of 128 C-bricks or 512 processors.


SGI Origin 3800 Technical Specifications

A multiple 39-U rack configuration with additional racks for I/O and disks as required:

Description	Minimum	Maximum
39U high racks	One compute rack and one I/O rack	16 Compute racks 8 I/O racks Disk racks as needed
Memory size	2GB (four C-bricks each with one 512MB bank installed)	1TB (128 C-bricks, each with four 2GB banks of memory installed)
Processors	Four C-bricks (16 processors)	128 C-bricks (512 processors) 8 C-bricks maximum per rack
Input/output	One I-brick, one P-brick standard	64 I/O bricks 8 I/O bricks maximum per rack
R-bricks	Two 8-port R-bricks	32 8-port R-bricks, 12 Metarouters
Power bays	One power bay per rack	Two power bays per compute rack One power bay per I/O rack

- Compute racks are configured with only C-bricks and R-bricks.
- I/O racks are configured with only D-bricks and I/O bricks.
- Nine D-bricks per I/O rack is the maximum (no power bays in rack)
- Processors are increased in increments that are based on the size of the system:
 - 16 processors to 128 processors in increments of 4, 16, or 32 processors
 - 128 processors to 512 processors in increments of only 32 processors
- Configurations that have more than 128 processors require metamemory DIMMs, which include directory memory
- The L2 system controller is mandatory for all compute racks
- The L3 system controller is optional

SGI Onyx 3800



Except for the addition of one or more graphics subsystems, the SGI Onyx 3800 entry configuration is identical to SGI Origin 3800. The base configuration includes two R-bricks and four C-bricks in the compute rack, one I-brick and one P-brick in the second (I/O) rack, and one G-brick housed in a third graphics rack. In addition, for each InfiniteReality2 graphics subsystem, the customer will also receive one keyboard and mouse, one audio card, and one USB extender. SGI Onyx 3800 systems can support up to 16 graphics pipelines. Each graphics rack supports two G-bricks, each with two graphics pipes, for a maximum of four graphics pipes per graphics rack. Like SGI Origin 3800, SGI Onyx 3800 supports up to 512 processors.

SGI Onyx 3800 Technical Specifications

A multiple 39-U rack configuration with additional racks for I/O and disks as required:


Description	Minimum	Maximum
39U high racks	One compute rack, one I/O rack, and one graphics rack	16 compute racks, 8 I/O racks 4 Graphics racks, Disk racks as needed
Memory size	2GB (four C-bricks each with one 512MB bank installed)	1TB (128 C-bricks, each with four 2GB banks of memory installed)
Processors	Four C-bricks (16 processors)	64 C-bricks (512 processors) 8 C-bricks maximum per rack
Graphics module	One G-brick (pipe not included)	8 G-bricks
Graphics pipe	One InfiniteReality2 or 3 pipe, one RM, one DG, one monitor	16 graphics pipes
Graphics options	Keyboard, mouse, audio card, USB extender	16 sets of options, one for each graphics pipe
Input/output	One I-brick, one P-brick standard	128 I/O bricks 8 I/O bricks maximum per rack
R-bricks	Two 8-port R-bricks	32 8-port R-bricks, 12 Metarouters
Power bays	One power bay per rack	Two Power Bays per compute rack One Power Bay per I/O rack

- Compute racks are configured with only C-bricks and R-bricks
- I/O racks are configured with only D-bricks and I/O bricks
- Graphics racks are configured with only G-bricks, 2 per rack
- Nine D-bricks per I/O rack is the maximum (no power bays in rack).
- Processors are increased in increments that are based on the size of the system:
 - 16 processors to 128 processor in increments of 4, 16, or 32 processors
 - 128 processors to 512 processor in increments of only 32 processors

- Configurations that have more than 128 processors require metamemory DIMMs, which include directory memor.
- The L2 system controller is mandatory for all compute racks
- The L3 system controller is optional

System Management Processor

An optional system-level controller based on a desktop workstation that runs Linux^(R) offers you the ability to manage large systems or clusters from a single administrative standpoint. Specialized system diagnostic software performed at this station allows SGI service personnel or customer system operators the ability to run complex diagnostics on the server. The ability to diagnose and predetermine system failures helps the customer avoid unexpected and costly downtime due to parts failure.



Benefits of SGI NUMA 3 Advanced Memory Design

Memory latency is critical to achieve maximum efficiency in a scalable architecture, and significant effort was expended on the SGI NUMA 3 architecture to drive memory latencies extremely low. The non-uniform memory access in SGI NUMA 3 is significantly lower than other NUMA implementations in the industry and also better than many Uniform Memory Access (UMA) architectures. Not only are local and remote access times significantly improved over previous generations of NUMA implementations from SGI, the ratio of remote to local memory is much closer, representing a memory performance very close to what would be expected from an UMA design. The ratio of remote to local memory access in the SGI 3000 family is 2 to 1, whereby the time it takes a processor to access the most remote memory location in the largest configuration (512-processor configuration) is only twice the amount of time it takes to access local memory in the smallest configuration. Memory access times for varying system sizes are listed in the table below for comparison.

SGI 3000 Series Memory Latency

# CPUs	Router Hops Max	Router Hops Average	Local Memory Latency	Worst-Case Remote Latency	Average Latency
16	1	0.75	175	285	257.5
32	2	1.38	175	335	296.3
64	2	1.69	175	335	315.6
128	4	2.47	175	435	356.6
256	5	3.48	175	485	408.3
512	7	4.74	175	585	471.6

Partitioning a System for Resiliency

The SGI 3000 family of systems has an innate ability to deliver high levels of resiliency and availability. The very flexible and modular SGI NUMA 3 architecture is designed with reliability, availability, and serviceability in mind, which customers can use to define a server solution that closely matches their application requirements. With the introduction of SGI NUMA 3, servers from SGI can now be partitioned into separate nodes to create an application environment that emulates a "cluster in a box."

Partitioning is defined as an ability to take a single distributed shared-memory (DSM) system, any model of the SGI 3000 family for instance, and divide it into a collection of smaller systems. The two primary characteristics of partitioning are:

- The ability to run individual partitions, whereby each partition runs in its own protected memory space with its own operating system kernel and behaves as a distinct, standalone system. Partitions can be booted, powered up or down, and rebooted without affecting the normal operation of the other partitions in the system.
- The partitions are tightly coupled, through the use of the system's interconnection network (NUMALink), as a low-latency, high-bandwidth interconnect. A failure that causes a kernel in one partition to panic will not cause a kernel in another partition to crash.

Partitioning can be thought of as a tightly coupled cluster that uses the lowlatency and high-bandwidth NUMALink Interconnect instead of a low-performance networking interface for the interconnect. The higher level of performance that results is directly related to the use of NUMALink to deliver information between partitions.

Fault isolation is one of the major reasons for partitioning a system. A software or hardware failure in one partition should be isolated from other partitions so that application availability in those other partitions is not affected. To accomplish this, a set of conventions is used to help define partitioning of systems by using IRIX 6.5:

- The minimum configuration for a partition is the combination of one C-brick, one I-brick, and possibly a separate power supply setup. Assuming all C-bricks are fully populated, the minimum partition size would be 4 processors and therefore is the minimum level of hardware isolation.
- Each partition must have the infrastructure to run as a standalone system. This infrastructure includes a system disk and console connection, both located in the I-brick.
- I/O bricks belong to the partition that the attached C-brick belongs to. If an I/O brick is dual-ported to two separate C-bricks, both of these C-bricks must be in the same partition. I/O bricks cannot be shared by two partitions.
- To allow communication to be independent, all intrapartition communication must be through a route that is contained within the partition.
- When the full system is greater than 64 C-bricks (256P), the minimum partition size is four C-bricks (16P, all C-bricks connected to a single R-brick).

IRIX Advanced Cluster Environment (ACE)

As some application environments benefit from a workflow that involves many small shared-memory servers, the ability to manage multiple servers in a simple but effective manner is the key to enhanced productivity. The IRIX Advanced Cluster Environment (ACE) line of products is specifically designed to help system managers improve work efficiency by simplifying the effort to manage many servers, both as individual nodes in a cluster and as multiple partitions in a larger shared-memory server. Used in cluster accounts for years, the tools in ACE for IRIX 6.5 are fully supported for production environments and deliver a single administrative image and single system view of a cluster or a multipartitioned shared-memory system. ACE for IRIX provides many features to effectively manage resources in a production environment:

-
- Message Passing Toolkit (MPT) provides a method to program scalable computer systems and arrays of workstations and servers. It includes the MPI development environment.
 - Job management is supported through the use of Miser, a batch job scheduling facility that balances batch and interactive CPU and memory usage in a single node. The optional Load Sharing Facility (LSF) tool is used to assign effective workload distribution and job scheduling across multiple nodes.
 - To manage a set of processes that run on different nodes in a cluster, Array Services provides a way to service conceptually related processes as a single "job."
 - To quickly identify and eliminate performance bottlenecks in a cluster, Performance Co-Pilot™ is a performance monitoring tool that supports monitoring across multiple nodes. Customers can choose from a number of add-ons for specialized monitoring and add more collector and monitor agents across the cluster.
 - The SGI management processor provides a centralized control point for system operators to manage and monitor multiple nodes or partitions, and logs their activity.
 - Software revision control and replication services can be managed with RoboInst, a tool that automates the process of installing operating systems, patches, and applications across multiple nodes or partitions.
 - Enlighten DSM is used by system operators for administration management and monitoring of applications, which enables users to immediately begin managing their environment at a workgroup level.
 - Clustered XFS™ (CXFS™) is a full-featured clustered file system that enables unparalleled data access across shared disks in a clustered configuration. CXFS is built upon XFS, a 64-bit journaled file system that delivers the industry-leading performance required for big data environments.

(c)2000 Silicon Graphics, Inc. All rights reserved. Specifications subject to change without notice. Silicon Graphics, Onyx, InfiniteReality, IRIX, Onyx2, and Geometry Engine are registered trademarks, and SGI Origin, InfiniteReality2, InfiniteReality3, NUMAflex, NUMALink, Performance Co-Pilot, XFS, CXFS, OpenMP, and the SGI logo are trademarks of Silicon Graphics, Inc. MIPS is a registered trademark. Linux is a registered trademark of Linux Torvalds. All other trademarks mentioned herein are the property of their respective owners.

The successful deployment of today's high-performance computing solutions is often obstructed by architectural bottlenecks. To enable organizations to adapt their computing assets rapidly to new application environments, SGI™ servers have provided revolutionary architectural flexibility for more than a decade. Now, the pioneer of shared-memory parallel processing systems has delivered a radical breakthrough in flexibility, resiliency, and investment protection: the SGI Origin 3000 series of modular high-performance servers.

Design Your System to Precisely Match Your Application Requirements

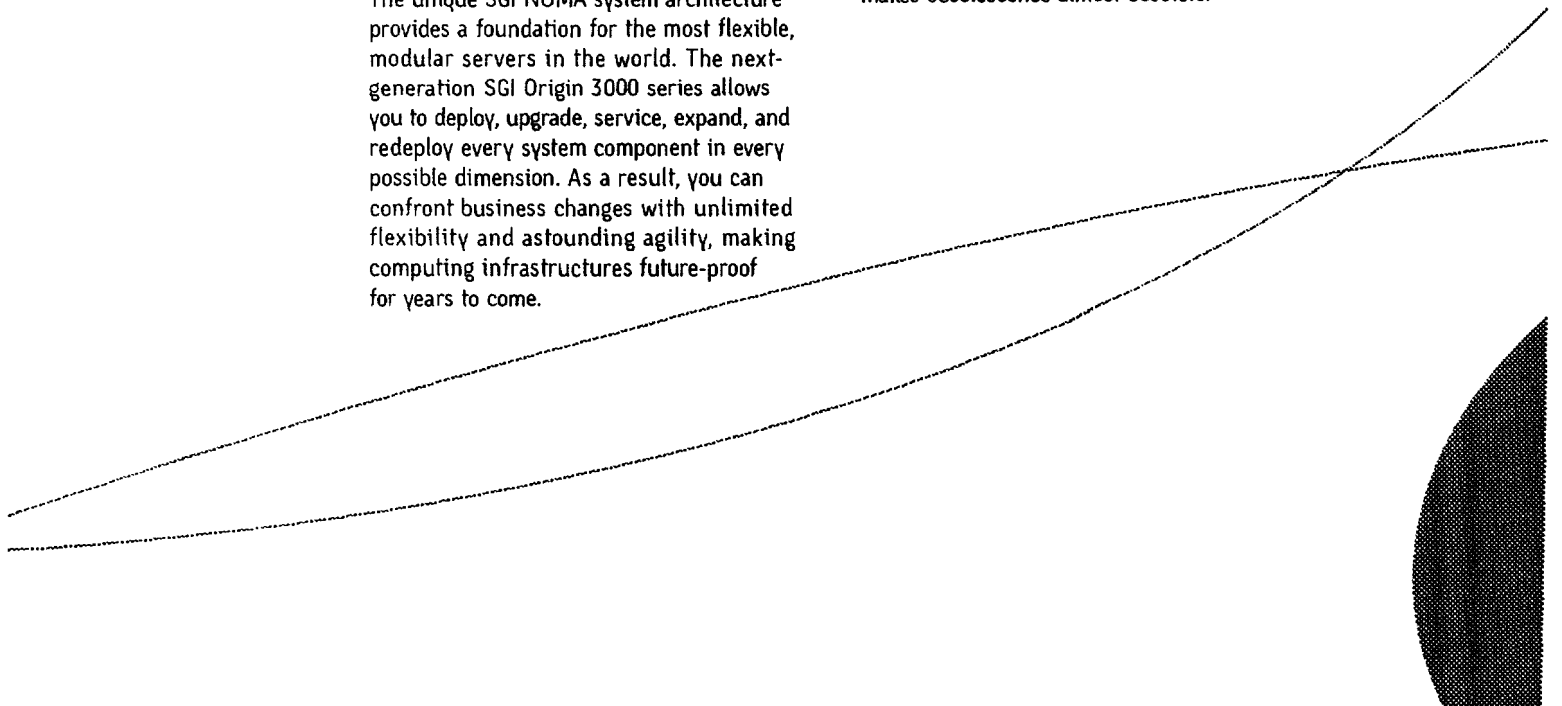
SGI Origin 3000 series systems take modularity to the next level. Building on the same modular architecture of award-winning SGI™ 2000 series servers, the SGI Origin 3000 series now provides the flexibility to scale CPU and memory, storage, and I/O components independently within the system. You can design a system down to the level of individual components to meet your exact application requirements—and easily and cost effectively make changes as desired.

Unmatched Flexibility


The unique SGI NUMA system architecture provides a foundation for the most flexible, modular servers in the world. The next-generation SGI Origin 3000 series allows you to deploy, upgrade, service, expand, and redeploy every system component in every possible dimension. As a result, you can confront business changes with unlimited flexibility and astounding agility, making computing infrastructures future-proof for years to come.

A New Snap-Together Approach

This new approach to server architecture allows you to configure—and reconfigure—systems brick by brick. Upgrade CPUs selectively to keep apace of innovation. Isolate and service I/O interfaces on the fly. Pay only for the computation, data processing, visualization, or communication muscle you need. Achieve all this while seamlessly fitting systems into your IT environment with an industry-standard form factor. Brick-by-brick modularity lets you build and maintain your system optimally, with a level of flexibility that makes obsolescence almost obsolete.




SGI Origin 3200



Two to Eight CPUs

This affordable system scales from two to eight processors in a shared-memory image without routers and can be clustered to leverage the power of hundreds of CPUs managed from a single point of administration. Brick-by-brick expandability makes it an ideal base for high-availability file, Web, or storage serving.


SGI Origin 3400



4 to 32 CPUs

A price/performance powerhouse, SGI Origin 3400 has two six-port routers to accommodate scaling from 4 to 32 CPU processors in a shared-memory image. For maximum flexibility, the processor and I/O bricks can be configured within the same or separate racks, allowing for extensive I/O expansion.

SGI Origin 3800



16 to 512 CPUs

Achieve unmatched power and unlimited scalability with the SGI Origin 3800 system. With the largest single-kernel, shared-memory image available, SGI Origin 3800 scales with snap-together modularity from 16 to 512 processors. Built-in Metarouters allow clustering to tens of thousands of CPUs, making supercomputing more accessible than ever.

Introducing NUMAflex—the new SGI Origin 3800 system. It's the most powerful, scalable, and flexible shared-memory system ever. It's the only system that can scale from 16 to 512 processors. It's the only system that can be configured to run a single-kernel, shared-memory image. It's the only system that can be configured to run a multi-kernel, shared-memory image. It's the only system that can be configured to run a multi-kernel, shared-memory image. It's the only system that can be configured to run a multi-kernel, shared-memory image.


The SGI Origin 3800 system was designed to make it easy to scale your system. It's the only system that can scale from 16 to 512 processors. It's the only system that can be configured to run a single-kernel, shared-memory image. It's the only system that can be configured to run a multi-kernel, shared-memory image. It's the only system that can be configured to run a multi-kernel, shared-memory image.

NUMAflex

Traditionally, scalability has charted a unilinear path. But ideally, scalability should follow the path that best meets the requirements of different applications—a path that you define.

Get the Configuration You Need

The SGI Origin 3000 series allows you to do just that, whether you need heavy-duty computation for weather simulation, massive storage for archiving bioinformatic data, high-performance I/O for media streaming, or integrated visualization of large data sets.



PCI Expansion

Hot-plug PCI maximizes the availability of your data. The patented SGI carrier allows easy access and upgrade or replacement.

NUMAflex


NUMAflex memory architecture

For optimum performance, NUMAflex memory architecture distributes data across the system's memory banks. This allows for a more efficient use of memory and a more scalable system. NUMAflex memory architecture is designed to support a wide range of applications, from high-performance computing to data-intensive applications. It provides a flexible and scalable memory architecture that can be configured to meet the needs of your application.

NUMAflex memory architecture is designed to support a wide range of applications, from high-performance computing to data-intensive applications. It provides a flexible and scalable memory architecture that can be configured to meet the needs of your application. NUMAflex memory architecture is designed to support a wide range of applications, from high-performance computing to data-intensive applications. It provides a flexible and scalable memory architecture that can be configured to meet the needs of your application.

Build and Maintain Your System—One Brick at a Time

To scale system performance to meet your needs, choose the bricks your applications require:



R-brick

Router Interconnect

As the structural building block of the system, the R-brick replaces the system bus; it's a high-speed crossbar connecting processors and memory and enabling each system component to be serviced or upgraded individually. Add infrastructure as you need it—from routerless desktide systems to an eight-port router multitrack configuration that delivers 512 processors in a single shared-memory environment.

C-brick

CPU Module

The basic C-brick module contains four MIPS® CPUs and local memory. A single crossbar memory controller delivers 200% greater CPU-to-memory bandwidth than previous generations. Now with four CPUs in a C-brick, the system offers a two-fold increase in CPU density, improving memory latency by up to 50% and minimizing the use of valuable floor space.

I-brick

Base I/O Module

The I-brick, standard in all systems, provides base I/O in a module and includes the system disk, CD-ROM, Ethernet, and four available PCI slots. As the system grows, customers have the option to partition the system for greater availability, using additional I-bricks as base I/O for each partition.

P-brick

PCI Expansion

For PCI expansion, a P-brick provides 12 hot-swappable PCI slots distributed over six 64-bit/66 MHz PCI buses. Total peak I/O bandwidth exceeds 3GB per second.

X-brick

XIO Expansion

For high-performance I/O expansion, X-bricks deliver four XIO slots that support HIPPI, GSN, VME, and digital video.

G-brick


Graphics Expansion

Tightly integrated InfiniteReality3 graphics add large-scale visualization capabilities for accelerated insight into complex data sets.


D-brick

Disk Storage

D-bricks provide modular JBOD mass storage for data-intensive applications. D-bricks support up to 12 drives, have dual-power supplies standard, and support drive capacities of 18GB, 36GB, and 73GB.



The SGI Origin 3000 series uses the latest expression of the revolutionary NUMA (nonuniform memory access) architecture. As today's preferred architecture for high-performance, multiprocessor systems, NUMA facilitates access to remote system resources with unmatched efficiency. It's the only way to effectively take advantage of today's high-powered CPUs. Unlike bus-based SMPs, the SGI NUMA architecture allows systems to increase shared memory to meet the growing CPU-to-memory bandwidth demands of additional processors. As a result, memory bandwidth grows proportionately as CPUs are added, making SGI NUMA systems inherently more scalable. It is the unprecedented modularity that makes the superior scalability of the SGI Origin 3000 series possible, allowing you to build large, high-performance systems brick by brick from the same components used to build an entry-level desk-side unit.




Breakthrough Architecture Enables a New Way of Computing


As the NUMA innovator, SGI has been years ahead of the industry since the introduction of the award-winning SGI Origin servers. Today, IDC calls NUMA "the architecture of the future." But the architecture of the future can be had today in a third-generation implementation at its most robust: the SGI Origin 3000 series.

Performance, Reliability, and Versatility
With their high bandwidth, superior scalability, and efficient distribution of resources, SGI servers are more than just highly modular. They are performance leaders. And now with the brick-by-brick flexibility of the SGI Origin 3000 series, performance gains come even faster—no matter how large the system gets.

• **Connectivity.** The SGI Origin 3000 series provides peak bandwidth for high-speed peripheral connectivity across all leading I/O solutions. And with bottleneck-busting speed, it supports the latest high-performance networking protocols.

• **Storage.** With ample expansion capabilities to accommodate today's ever-growing data sets, the most advanced storage technologies are supported—from 100MB-per-second Fibre Channel to 40MB-per-second Ultra SCSI. Protect your data with RAID storage solutions from SGI, offering support for RAID levels 0, 1, 0+1, 3, and 5, plus global hot sparing across all disk drives. High-end HSM solutions are available for enterprise-wide storage requirements.





• **Availability.** To provide the reliability today's applications demand, the SGI Origin 3000 series can deliver uninterrupted availability through ECC memory, redundant power and cooling, and hot-pluggable disks and PCI, along with IRIS FailSafe™ software.

• **Graphics.** For integrated graphics, the SGI Origin 3000 series supports multiple InfiniteReality3 graphics subsystems, enabling high-performance visualization of large data sets.

A Natural Extension That Protects Your Investment

As the newest generation of SGI servers, the SGI Origin 3000 series allows you to reap new flexibility and performance benefits without sacrificing what has come before. It will leverage and extend your investments in existing shared-memory applications. And since the series uses the powerful UNIX® system-based, 64-bit IRIX® 6.5 operating system, it is fully compatible with other SGI workstations and servers. With the same tools and a familiar operating system, you can integrate the SGI Origin 3000 series without any retraining. Your investment in application software and training is protected, and the availability of open systems software for the future is ensured—facilitating the ongoing development of new applications.

A Fully Supported Solution

The unprecedented flexibility of the SGI Origin 3000 series provides endless possibilities for designing custom solutions for today's competitive-edge applications. To architect and deploy solutions that capitalize on the strengths of the SGI Origin 3000 series, turn to SGI Professional Services. Industries as varied as manufacturing, media, energy, finance, and science rely on SGI consultants for complex data visualization, high-performance data mining and warehousing, e-commerce, and Internet solutions, among others. It's powerful support that leverages resources and maximizes return on investment.

A Breakthrough System for the 21st Century

With the highest level of flexibility, unlimited scalability, and stronger performance than ever before, the new SGI Origin 3000 series delivers an unequalled combination of flexibility and power to meet the changing computing requirements of any organization—today and for years to come.

NUMAflex